SAID: A Control Protocol for Scalable and Adaptive Information Dissemination in ICN

Jiachen Chen*, Mayutan Arumaithurai*, Xiaoming Fu*, and K.K. Ramakrishnan†
*University of Göttingen, Germany. Email: {jiachen,arumaithurai,fu}@cs.uni-goettingen.de
†University of California, Riverside, USA. Email: kk@cs.ucr.edu

arXiv:1510.08530v1 [cs.NI] 29 Oct 2015

***Abstract*—Information dissemination applications (video, news, social media, *etc.*) with large number of receivers need to be efficient but also have limited loss tolerance. The new Information-Centric Networks (ICN) paradigm offers an alternative approach for reliably delivering data by naming content and exploiting data available at any intermediate point (*e.g.*, caches). However, receivers are often heterogeneous, with widely varying receive rates. When using existing ICN congestion control mechanisms with in-sequence delivery, a particularly thorny problem of receivers going *out-of-sync* results in inefficiency and unfairness with heterogeneous receivers. We argue that separating reliability from congestion control leads to more scalable, efficient and fair data dissemination, and propose SAID, a Control Protocol for Scalable and Adaptive Information Dissemination in ICN. To maximize the amount of data transmitted at the first attempt, receivers request *any next packet* (ANP) of a flow instead of next-in-sequence packet, independent of the provider's transmit rate. This allows providers to transmit at an application-efficient rate, without being limited by the slower receivers. SAID ensures reliable delivery to all receivers eventually, by cooperative repair, while preserving privacy without unduly trusting other receivers.**

I. Introduction

Large scale information dissemination applications like video streaming (YouTube, NetfFlix, *etc.*), online social networks (Facebook, Twitter, *etc.*) and news/entertainment (CNN, BBC, RSS feeds, *etc.*) have become common. Many of these applications have limited loss tolerance and depend on the network to provide efficient, fair and reliable content distribution. While IP multicast was designed for large-scale information dissemination, the inability to have an effective congestion control solution, especially in the presence of heterogeneous capacity to receivers, has been a limitation. The advent of Information-Centric Networks (ICN) offers us an opportunity to take a fresh look at the potential solution approaches.

A key goal for publishers sending data at an application-efficient rate across the entire receiver population is to be not limited by the slower receivers. Due to the absence of a network layer mechanism to control the delivery rate at the receiver end, previous solutions have either sought to push all the data onto the path, overlooking congestion and unfairness and using end-end unicast recovery, or seek to slow down the sending rate to the slowest receiver [1]. Alternatively, the use of unicast, with the associated inefficiencies, has become the norm. The use of multicast at the application layer (*e.g.*, SCRIBE [2]) exploits TCP's congestion control mechanisms for ensuring loss-free delivery on an end-end basis. However, application layer solutions are network topology unaware and achieve lower efficiency (caused by the end-hosts replicating the packets rather than the routers) than what an effective multicast solution could be expected to achieve.

Information-Centric Networking (ICN) is a new paradigm with the potential to achieve efficient large scale information delivery. Content-Centric Networks (CCN [3]) or Named Data Networking (NDN [4])[1] is a representative ICN approach. The shift from a "location-focused" network to a "content-centric" network allows more efficient data sharing, thanks to name-based forwarding. With in-network caches, NDN can exploit temporal locality among the consumers to outperform existing IP network-based approaches for a variety of information delivery situations. Moreover, since entities in the network exchange information primarily on content names, the identities of the consumers need not be revealed, thus preserving user privacy. Additionally, data integrity and provenance can be established based on the per-packet data generator-signature required in NDN.

Although NDN does not mandate a congestion control mechanism, most of the proposed solutions [5]–[9] choose to use a TCP-like receiver-driven mechanism to limit the number of requests (window) outstanding from a receiver. Since the network maintains flow balance, where one Interest retrieves at most one Data packet, these mechanisms adapt the window using the Additive Increase Multiplicative Decrease (AIMD) principle, much like TCP [10], [11]. Compared to a sender-driven rate control approach, such a receiver-driven approach has the benefit of the consumer being able to control the receive rate. When considering efficient large scale data dissemination where every piece of data is consumed by a large number of receivers, TCP-like mechanisms for receiver-driven multicast can have significant shortcomings, especially with the path to the receivers having different capacities. With such heterogeneity, a problem we observe is that of receivers being **out-of-sync** *even with optimal policies for managing the cache*. The out-of-sync problem can be briefly described as follows: NDN routers cache content as they are forwarded. When there is temporal locality of requests from different receivers, a router that has the cached content can respond. However, with receiver heterogeneity, the requests from receivers even for the same data items can diverge over time. Requests from the faster receivers can be well ahead of those from slower receivers. Eventually, when the gap between the

[1]In this work, we do not distinguish between NDN and CCN since we only refer to the fundamental communication model.

faster and slower receivers becomes too large, their requests can no longer be aggregated at the intermediate routers or satisfied by the cache. The slower receivers' requests will have to be satisfied by the content provider (via retransmissions) as a separate flow. These retransmissions will compete for the bottleneck bandwidth and the overall throughput can therefore dramatically reduce. This is a fundamental issue as long as there are heterogeneous receivers and routers with limited cache sizes. The problem can be exacerbated with scale, thus occurring even more often in the core of the network.

In this paper, we propose a control protocol for Scalable and Adaptive Information Dissemination (SAID) in ICN, a novel mechanism enabling large scale efficient data dissemination. We leverage the receiver-driven framework of NDN with enhancements to overcome the out-of-sync problem. SAID achieves efficiency by separating reliability from congestion control. By requesting for "any" packet that comes next, instead of the "next packet in sequence", SAID is able to maximize the amount of packets delivered *at the first attempt*. Reliability is achieved by consumers performing information-centric repair which also has the attractive feature of preserving privacy and data integrity. While our design is framed in the context of NDN, we believe SAID can also be used by other reliable multicast solutions. The contributions of this paper include:

- An analytical and emulation-based study on the out-of-sync problem that shows it will occur in real networks even with large in-network caches as long as there is receiver heterogeneity. This results in lower useful throughput, higher publisher and network load (see §II);
- A new reliable multicast framework that seeks to first maximize the useful throughput, by consumers requesting for *Any Next Packet* (ANP Delivery, see §IV-A). Reliability is then achieved via a repair mechanism that leverages NDN's capability for receivers to request a data item from any network node including other receivers, while preserving privacy and data integrity (see §IV-B);
- A receiver-driven congestion control mechanism tailored for the ANP Delivery that enables *each* receiver to obtain its fair share of the bottleneck link while maintaining an application-efficient sending rate that doesn't necessarily slow down to the slowest receiver's rate (see §V);
- Evaluations on our prototype show that SAID is able to outperform ICP (by 50%) and pgmcc (by almost 100%) in terms of aggregate throughput without sacrificing network fairness; in a file content delivery application, SAID can reduce aggregate network load (by up to ∼46%) and transmission completion times (by more than 50%) compared to ICP. Compared to pgmcc, SAID reduces completion times by 40% while only increasing network load by 10% (see §VI).

II. Problem with Existing Congestion Control Mechanisms – Out-of-Sync

We first study the out-of-sync problem in CCN and show how it reduces the benefit of in-network caches and the use of pending interests. We demonstrate this with an emulation

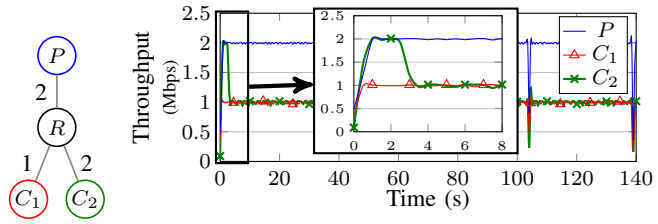

(a) Topology (b) Throughput on the endhosts

Fig. 1: Out-of-sync problem emulated in a simple topology.

using CCNx 0.8.0 along with a congestion control mechanism similar to ICP [5]. We show through an analytical model that our observations on the out-of-sync problem is systemic with heterogeneous receivers, and should be expected even with other receiver-driven, in-sequence congestion control mechanisms.

A. Demonstration of out-of-sync in an emulated scenario

To clearly demonstrate the out-of-sync problem and its cause, we use a simple emulation performed in Mini-CCNx. The network topology and the link rates (in $Mbps$) are shown in Fig. 1a. The latency on all the links is $2ms$. Router R has a 50 packets cache[2]. Consumers C_1 and C_2 start to request the same content ($\sim 35MB$ in size, $8,965pkts$) from provider P at the same time. The throughput of the end hosts are shown in Fig. 1b. For the first 2.5 seconds, the PIT and the network cache benefit both receivers. Requests from C_1 either get aggregated or get a cache hit at R. Overall network throughput in this period is $3Mbps$ (sum of downstream link capacities of R). Subsequently, with heterogeneous receiver rates, the receivers' requests deviate farther apart. The requests from C_1 can no longer be satisfied by the cache and they are forwarded to P, to be treated as a distinct flow. The response to these requests start to compete for the bandwidth on the link from P to R and the receive rate of C_2 is thus affected. Since the congestion control protocol tries to achieve fairness between the receivers (flows), the receive rate of the two consumers becomes $1Mbps$ each, and the overall network throughput reduces to $2Mbps$, thus under-utilizing the bandwidth by 33%. For the entire transfer, we observed that $< 2\%$ of the requests from C_1 see a cache hit.

Similar results occur in more complex topologies (*e.g.*, Fig. 10 without dotted links). The consumers C_{11}, C_{21}, C_{31} and C_{41} start to request the same content from provider P_1 at the same time (C_{51} is not active in this emulation). The cache size is $100pkts$ and the content size is $10,000$ packets. The receive rates at the consumers are shown in Fig. 2a and the # of requests (transmissions) per packet observed by P_1 is shown in Fig. 2b. We can see that, similar to the pervious more simple case, the consumers once again get out-of-sync soon after the transmission sequence starts. The # of transmissions increase and the receive rate drops. For the first

[2]We use a relatively small cache size to quickly demonstrate the out-of-sync problem. However, this is fundamental and occurs even with much bigger caches. Please see §II-B for the relationship between the cache size and the heterogeneity allowed among the receivers.

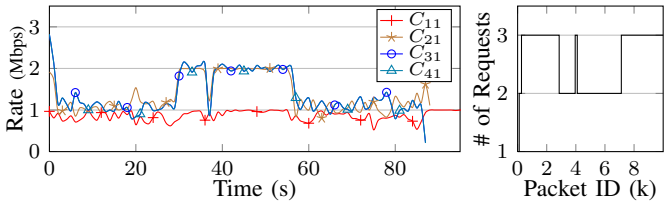

(a) Receive Rate on consumers (b) # of req./pkt

Fig. 2: Out-of-sync: On larger dissemination tree.

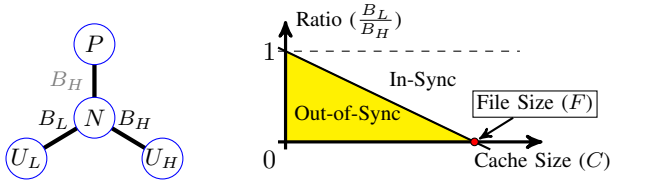

(a) Single-branching model (b) Relationship between out-of-sync and request rate, cache size and flow size

Fig. 3: Out-of-sync in a single-branching model.

30 seconds, the aggregate throughput is only around $4Mbps$ (the ideal throughput is $9Mbps$). During the intervals 30–36 and 40–55 seconds, the faster receivers get in-sync again, due to the randomness in the network and cache occupancy. The packet sequences between $2,800$–$3,950$ and $4,100$–$7,100$ are transmitted twice and the throughput of C_{21} through C_{41} increases during these two time periods. But even during this time, the aggregate throughput can only reach $7Mbps$ rather than the $9Mbps$ achievable in the ideal case.

Through these emulations, we see that due to the heterogeneity of the receivers, the cache in the intermediate routers might not be enough to absorb the difference in the request rates of the fastest and slowest receivers. This is the occurrence of "out-of-sync" problem. When the slower receivers re-issue requests, these requests are seen as being for a different "new" flow, since they can no longer be aggregated or be satisfied from the cache at the routers. These "new" flows will then compete on the network links with packets of the original flow. In some cases, this would even be with faster receivers on the common links, and affect their download rate as well.

The out-of-sync problem can happen even when all the receivers start their requests for the sequence of data packets at the same time, even with the optimal cache replacement policy. Note that the provider has to re-transmit packets as long as the intermediate router drops the packets within the gap (the difference in the sequence number of the packet requested by the fastest and slowest receiver). When the gap is larger than the available cache size for the flow, no matter which packet the replacement policy chooses, an additional transmission from the provider is required.

B. Analytical model for out-of-sync occurrence

Receiver-driven feedback-based in-sequence congestion control protocols (*e.g.*, TCP) share the following features: *1*) each data consumer has a local view of the request as if he is the only consumer in the network, *2*) all the data consumers tend to get a (statistically) fair-share of bandwidth, and *3*) the packets in a data object are requested in-sequence and out-of-order is seen as an indication of congestion. We realize that almost all the existing congestion control protocol proposed for CCN fall into this category.

To show the universality of the problem, we generalize the model for congestion control by assuming a best-case scenario where each receiver is receiving a flow of data with a constant bit rate which is exactly the fair-share that receiver can get. We analyze the maximum heterogeneity that can be supported given a certain cache and flow size while the receivers still remain in-sync till the end of the flow. Since we are focusing on a single flow, we also assume a simpler case that the network status (*i.e.*, available bandwidth, cache size and latency) does not change during the lifetime of the flow.

We start with a more precise definition of out-of-sync.

Definition II.1 (Out-of-sync)**.** *Consider a network with multiple routers interconnected in a tree topology and a flow f that has a provider on the top and receivers at the leaves of the tree. At a branching router N, where available bandwidth to the downstream receivers is in the range $[B_L, B_H]$, the out-of-sync occurs when the difference in the received file size (the gap, G) between the fastest and slowest receiver in the sub-tree below N is larger than the available cache size C for flow f.*

We then use a single-branch model to demonstrate the relationship among the cache size, flow size and receiver heterogeneity in a simple scenario.

Lemma II.1. *For a branching router N with cache size C in a dissemination tree, with the request rates of the immediate downstream links in range $[B_L, B_H]$, to avoid out-of-sync, the following condition should hold,* i.e.*:*

$$\frac{B_L}{B_H} \geq (1 - \frac{C}{F}), \tag{1}$$

where F is the size of the flow.

Proof. According to Def. II.1, to avoid out-of-sync at N, it is sufficient to consider 2 immediate downstream links with the largest and smallest available bandwidth. That is, we can consider the single-branching topology in Fig. 3a such that the two data consumers (U_L and U_H) are requesting for a same flow with size F and their available bandwidth are B_L and B_H ($B_L \leq B_H$).

When the receivers are in-sync, the request sent upstream by N targets a downstream rate of B_H, matching the receive rate of the faster receiver (this is true for all protocols where the network node does not perform an explicit congestion control function and depends on the receivers to generate an appropriate request rate). The download period for U_H is:

$$t = \frac{F}{B_H}.$$

The maximum gap G between the U_L and U_H is therefore:

$$G = (B_H - B_L) \times t = (1 - \frac{B_L}{B_H}) \times F$$

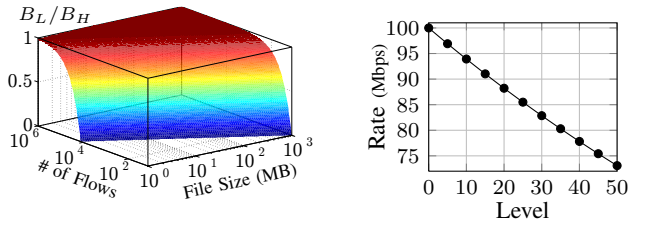

(a) Relationship among out-of-sync criteria in real router (b) Minimum allowed download rate *vs.* hierarchical level

Fig. 4: Difficulty of keeping heterogeneous receivers in-sync.

According to Def. II.1, to keep the consumers in-sync, we need:

$$G = (1 - \frac{B_L}{B_H}) \times F \leq C. \tag{2}$$

We can get (1) by reforming (2). □

The requirement for clients in-sync (equation 1) is presented in Fig. 3b. We can see that the requirement for being in-sync cannot be satisfied when the heterogeneity ($\frac{B_H}{B_L}$) is larger, the flow size F is larger, and/or available cache size C is smaller.

Since the request and data paths in large scale data dissemination usually form a tree structure rooted at the data provider, we then study the in-sync requirements in a k level tree.

Theorem II.2. *For a dissemination tree with k levels and every intermediate router having cache size C, all the receivers will be in-sync only when the available bandwidth between the fastest receiver and the slowest receiver follow:*

$$\frac{B_L}{B_H} \geq (1 - \frac{C}{F})^k \tag{3}$$

Proof. We prove this theorem by contradiction. Suppose that the request rates of the highest and the lowest receivers satisfy

$$\frac{B_L}{B_H} < (1 - \frac{C}{F})^k, \tag{4}$$

and all the receivers are in-sync.

Without loss of generality, we assume that the receiver with the lowest rate B_L is a downstream consumer of a router N_t at level $t \in [1, k]$. Let $B_{H,t}$ be the highest request rate among the downstream consumers of the router N_t. Note that the consumer with request rate B_H does not have to be the immediate next hop of N_t, the intermediate routers will always forward requests according to the fastest receiver. According to Lemma II.1, we have

$$\frac{B_L}{B_{H,t}} \geq 1 - \frac{C}{F}. \tag{5}$$

According to (4) and (5), it follows that

$$B_{H,t} < B_H \times (1 - \frac{C}{F})^{k-1}. \tag{6}$$

The router N_t is a downstream consumer of a router N_{t-1} at level $t-1$. Similarly, let $B_{L,t-1}$, $B_{H,t-1}$ be the lowest and highest rates among the downstream consumers of the router N_{t-1} respectively. Since $B_{H,t} \geq B_{H,t-1} \geq B_{L,t-1}$, according to Lemma II.1, we have

$$\frac{B_{H,t}}{B_{H,t-1}} \geq \frac{B_{L,t-1}}{B_{H,t-1}} \geq 1 - \frac{C}{F}. \tag{7}$$

According to (6) and (7), it follows that

$$B_{H,t-1} < B_H \times (1 - \frac{C}{F})^{k-2}.$$

By the similar argument, we can show that

$$B_{H,t-2} < B_H \times (1 - \frac{C}{F})^{k-3},$$
$$B_{H,t-3} < B_H \times (1 - \frac{C}{F})^{k-4},$$
$$\vdots$$
$$B_{H,1} < B_H \times (1 - \frac{C}{F})^{k-t} \leq B_H.$$

Since the highest rate of downstream consumers of the router at level 1 should be B_H, *i.e.*, $B_{H,1}$=B_H, we reach a contradiction and the proof is completed. □

From Theorem II.2 we can see that with the number of levels (k) increases, the gap between the fastest and the slowest receiver can become larger. The design of the in-network cache helps in absorbing the heterogeneity of the receivers.

Now we show that out-of-sync is difficult to avoid in a CCN router deployment.

Remark. *The problem of receivers going out-of-sync persists with receiver-driven feedback-based in-sequence congestion control protocols as long as there are heterogeneous receivers.*

The cache size at a router will inevitably be much smaller than the total amount of content available in the network. According to [12], CCN requires a 25TB cache for a 50% hit rate on Youtube data and 175TB cache for a 50% hit rate on BitTorrent data. But [13] suggests that a deployable CCN router (with ∼$1,500 overall hardware cost) would likely have around 100Gb of cache at current costs. Thus, a router cache will be much smaller than the required cache size for the kind of content accessed in current day networks.

For tractability, we assume that concurrent flows in a router share the 100 Gb cache size equally. The relationship between the bandwidth ratio ($\frac{B_L}{B_H}$), file size (F) and number of flows is shown in Fig. 4a. The intersection of the curve with XY plane is where C=F (# of flows=$\frac{100Gb}{F}$). The region behind the curve represents the region receivers are out-of-sync. Note that both the X- and Y-axis are log-scale, which means the area when the receivers are in-sync is just a very small portion of the overall region.

Although a core router might have relatively larger cache, the number of concurrent flows on that router is also correspondingly large. The available proportion of cache for each flow is therefore still quite small. If a set of receivers request 20M bytes of data through a core router with 100k concurrent flows, the ratio of rates should satisfy the following:

$$B_L/B_H \geq 1 - (100Gb/100k)/160Mb = 0.99375$$

If B_H can reach 100Mbps, B_L should be ≥99.3Mbps. When we apply the hierarchical tree model in Theorem II.2, the

TABLE I: Overview of reliable multicast protocols.

	Network Efficiency	Network Fairness	Application Efficiency	App.-specific Reliability
SRM	✓	✗	✓	✓
pgmcc	✓	✓	✗	✗
CoolStreaming	✗	✓	✓	✓
ICP	✗	✓	✗	✓

minimum download rate *vs.* level is plotted in Fig. 4b. Even if we have 50 levels in such a hierarchy, the minimum required download rate is still >73Mbps. This is difficult to achieve due to the number of flows that may be multiplexed on a given link.

C. Existing solutions for preventing out-of-sync

Existing NDN applications such as audio/video conferencing [14] seek to implement in-sync delivery at the application layer. Initially, the communicating parties negotiate the data delivery rate (*e.g.*, $n\ pkt/s$) to maintain real-time communication, skipping packets arriving late. Thus, instead of requesting the "next-in-sequence" packets, end hosts only request packets that can possibly arrive in time.

But this complicates the application, requiring it to negotiate the delivery rate a priori and also continuously monitor packet arrival and consumption rates which may vary over time. The content provider and consumer may find it difficult to set an appropriate sending rate given the network's changing available bandwidth. Moreover the user experience may be sacrificed, with "gaps" or interruptions - whether for conferencing or streaming content (*e.g.*, video-on-demand(VoD)). When the missing content has to be retrieved (*e.g.*, for VoD), this also results in out-of-sync behavior.

Thus, to maximize the utility of the network for file delivery and VoD, and to provide better support for real-time communications, there is a need for a transport-layer protocol that can keep the data receivers in-sync while retrieving data.

III. Protocol requirements

To motivate our design, we sum up the requirements for an efficient control protocol for large-scale data delivery and identify the issues with some selected, existing solutions. We categorize existing solutions and consider one representative (see Table I): SRM [15] for reliable IP multicast without congestion control; pgmcc [1] for reliable IP multicast that aligns to the slowest receiver; CoolStreaming [16] for application-layer multicast, and ICP [5] for a receiver-driven solution in ICN.

Network Efficiency: An efficient control protocol should maximize the utility of each packet sent from the data provider, seeking to deliver packets *on the first attempt* to as many receivers as possible – it is desirable keep the data consumers in-sync (where possible) when the data provider responds to their requests. Receivers being out-of-sync causes higher overall transmissions, resulting in lower useful network throughput, higher aggregate network load and provider load.

Network Fairness: Considering one-to-many flows, each flow should obtain a fair share of the bottleneck bandwith on *each path in the dissemination tree*. SRM injects data from the provider towards all the paths, ignoring this need to provide a fair share of the bottleneck capacity.

Application Efficiency: We argue that a data provider sending at the slowest receiver's consumption, like pgmcc is not application-efficient – a very slow receiver has an unduly significant influence on the application. An alternative is to enable slower receivers to seek alternate means (*e.g.*, get missing packets/repair from other faster receivers). Repairs certainly add to the load on some segments of the network, but the intent is to achieve a tradeoff between network load and session completion time.

Application-specific Reliability: Due to the difference in the requirements of the applications, The protocol should be general and flexible to provide variable levels of reliability based on application need. *E.g.*, file delivery needs reliability without in-order delivery; VoD would tolerate out-of-order delivery in a small range (within the play out buffer size) and limited packet loss; and conferencing applications do not need reliable delivery up to a point.

As we show in Table I, existing approaches (characterized by the representative set) do not meet all these requirements, thus motivating our work on SAID.

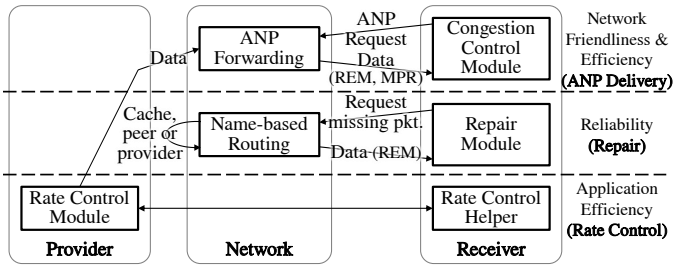

Fig. 5: SAID overview.

IV. SAID Framework

Because the need for reliability and in-order delivery varies across applications, we argue that the one-many (or many-many) information delivery and associated congestion control mechanisms in multicast and ICN/NDN do not need to enforce in-sequence delivery and reliability by default. Such solutions result in overheads (*e.g.*, receivers being out-of-sync, with NDN/ICP) or inefficiency (sending at the slowest receiver's rate, with pgmcc). Therefore, we propose SAID, a general framework for multiparty information delivery, that provides efficiency and flexible reliability in the ICN/NDN context. The principles in SAID can be generally applicable for multicast as well. Fig. 5 provides an overview of the proposed framework.

SAID is designed to tolerate missing packets while maximizing the delivery probability of the packets on the first attempt (for *network efficiency*), minimizing the need for retransmissions. This is difficult to achieve in NDN's sequence-specific request model with varying network conditions and heterogeneous receiver capacities. Therefore, receivers in

SAID use an ANP Delivery model that requests for "any-next" packet and let the network decide which packet can be delivered.

Reliability can then be achieved by receivers leveraging NDN's sequence-specific request model to recover missing packets without sacrificing privacy and trust, since the request only needs to identify the name of the missing packet and not necessarily the receiver's name or location (SAID's repair). Such repair requests can be satisfied by the content provider, caches or even other receivers. Applications can prioritize the packets to be recovered, or even skip this.

Many applications (*e.g.*, file delivery) have elastic delay requirements. The transmission rate of the provider chosen would be based on a tradeoff between the (network, provider) load and the content delivery completion time. A higher sending rate reduces the completion time for the faster receivers, but could also result in higher retransmissions to reliably deliver to the slower receivers. A lower sending rate may seek to operate at the opposite end of the spectrum – sacrificing completion time for lower network load. As a selective component at the application layer, SAID accommodates a mechanism for the provider and consumers to negotiate the provider's sending rate. SAID allows the provider to pick a sending rate that ensures a application-dependent proportion (*e.g.*, 60%) of the receivers can receive at that rate. An *ACKer* among the receivers is chosen to pace the flow, similar to pgmcc. But, unlike pgmcc, the ACKer does not have to be the slowest receiver. Although we do not cover the negotiation protocol for ACKer selection, due to space limitations (refer to our technical report [17] for a more detailed description), our evaluation of a file content delivery application in a real-world topology includes the ACKer selection protocol for controlling the sending rate (see §VI-B2).

A. ANP Delivery Model

To overcome heterogeneous receivers getting out-of-sync, and maximize the likelihood of delivery of a packet in the first attempt, we propose a slight modification to the existing communication model of receivers asking for a specific packet (as is currently done by NDN and even TCP, *etc.*), whereby receivers ask for *any* subsequent incoming packet.

We reuse the existing Interest packet in NDN for requests (*i.e.*, prefix, selector and nonce in TLV format) with an extra field indicating this request is seeking "any next Data" that can satisfy the prefix in the Interest header. In the CCNx based implementation we modify the `Exclude` field (to "exclude" all the previous Data available at the node).

On receiving ANP Delivery requests, an upstream router does not perform the usual check of the Content Store (cache) to avoid transmitting duplicate packets. Instead, it places the prefix, the incoming face (as in NDN) and a pending request counter (PR) into its PIT. PR is incremented on receiving an Interest packet with the same prefix and the routers propagate the maximum of the PR value upstream until this propagated value reaches the 1^{st} hop router of the data provider. This modification is more space-efficient than the current sequence-specific requests. For n requests of the same flow, SAID only needs 1 PIT entry and 1 counter compared to n separate PIT entries with the existing approach. When a router receives a Data packet of the flow (which satisfies the ANP request), it is forwarded on the outgoing face and PR is decreased by 1. The PIT entry will be removed when PR becomes 0.

With this subtle change, receivers in SAID can get fresh data without getting out-of-sync, thus enhancing network efficiency. This modification is also exploited for congestion control, elaborated in §V.

B. Achieving Reliability via Efficient Repair

For applications that need reliable delivery, we depend on the application interface to retrieve missing packets. Applications should be able to determine which packets are of higher priority. SAID leverages NDN's sequence-specific data retrieval capability to retrieve the missing packets from any network node including peers.

1) Repair among Consumers

Similar to traditional NDN, in addition to the data provider, all the end hosts in SAID that have (part of) the data can also support repair requests. But, the difficulty lies in identifying and routing the requests to the receivers that have these missing data packets. Moreover, since the repair packet has to share the bandwidth on the dissemination tree, independent of whether the repair request is delivered from a cache in the network or the provider, the repair has the potential to aggravate the congestion at a bottleneck link in the path from the provider to the receiver. Therefore, we propose a local FIB propagation optimization that can potentially mitigate the impact of repair on an already congested link.

With SAID, after receiving (part of) the data from the provider, receivers will flood the prefix of the data over a limited number of hops. *E.g.*, in Fig. 6, after receiving the packets, C_1 propagates prefix `/arXiv/SAID.pdf/_v1` over a 2 hop range. When C_2 requests for `/arXiv/SAID.pdf/_v1/_s20`, R_7 will forward the request to R_6 and eventually with get a response from C_1. The repair can then bypass the congested link R_5–R_7 and still increase the overall useful throughput to C_2. We believe multiple paths will continue to exist, just as we see in current networks [18]–[20].

Similar to peer repair solutions [15], [21]–[23], two major concerns with "peer receiver-based repair" are privacy for data consumers and the integrity of the data received from a peer [24]. In IP (location-based) networks, the data consumer in most cases will reveal his own location (identity) while requesting from a peer (except those that explicitly choose to remain anonymous, *e.g.*, by using Onion Routing). NDN, however give these applications a natural way to achieve privacy and trust by default. Since every data receiver requests repair via a Content Name, the receiver does not have to reveal his identity. Even for peers that perform the local FIB propagation, they propagate the prefix of the flow or the chunk (a chunk represents n packets) rather than their own identity. Every Data packet in NDN has a key digest of the data provider and a signature for the data content from the provider. The data consumer can easily check the integrity of the packet regardless of whether it is received through repair from other receivers or from the original data provider.

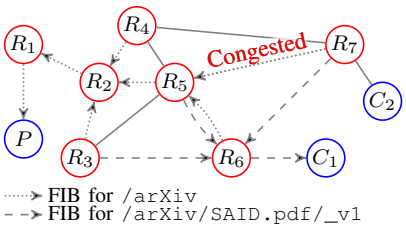

Fig. 6: Example of different repairs.

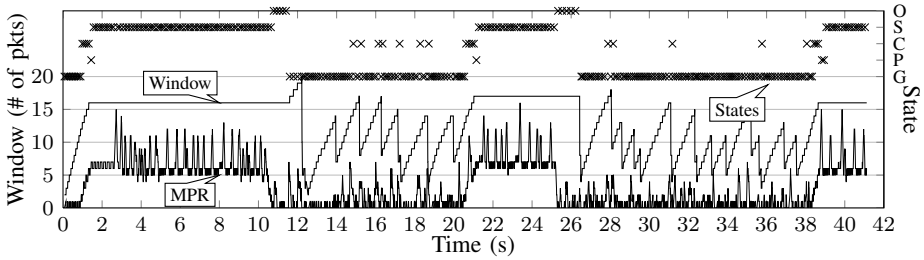

Fig. 7: Result of window control at the receiver using state machine.

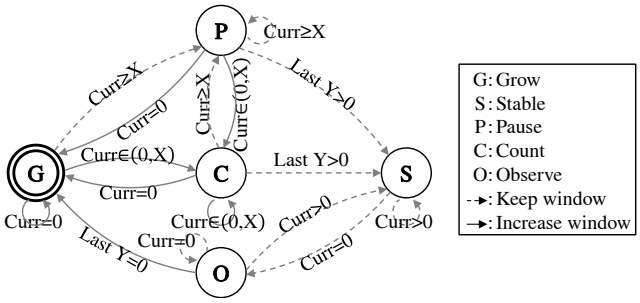

Fig. 8: State machine for receiver window increase.

2) Prefix Granularity

The simple mechanism we describe above may still be inefficient. A receiver, say C_1 would propagate a FIB entry when it receives one or more packets of a flow so that repair requests for slower receivers can be redirected to C_1. While propagating a FIB entry for each packet can result in excessive overhead, propagating a single entry for the entire content after receiving *all* the packets of the flow may not be useful since C_1 can only respond to repair requests at the end of the flow. To achieve a balance between the repair efficiency and reduced FIB size, we group n packets into a "chunk" and replace the `segmentID` in the ContentName with `chunkID/segmentID`. *E.g.*, if n=100, the name of the packet with segmentID=205 would be `/arXiv/SAID.pdf/_v1/_c2/_s205`. Every packet still has a globally unique ContentName. But C_1 can propagate `/arXiv/SAID.pdf/_v1/_c2` after receiving packets 200-299 for timely repair for other receivers. The FIB entries created will be much smaller as well.

V. Congestion Control for ANP Delivery

A pure network based congestion control solution for multicast that divide sthe outgoing link bandwidth fairly for each flow going to the *next hop* is unable to accommodate indications of reduced demand from flows that have a limit downstream, especially with heterogeneous receiver bandwidths. This results in inefficiency, as was suggested by [25], and we show it in §VI-A3 SAID chooses to use an end-system assisted solution (receiver-driven) to avoid such inefficiency. First, SAID retains the "flow balance" property suggested in NDN, in which every Interest packet will have at most one Data packet returned. Receivers then maintain a window of unsatisfied ANP requests, reflecting the maximum number of packets that can be in flight towards this receiver. We can adapt this window, just as TCP, with an AIMD approach, to find the capacity from sender to receiver. However, the sender's rate is not controlled by this window, when there are heterogeneous receivers. The difference between the sending and receiving rate naturally results in the receiver missing some of the packets in the sequence (we use the term "holes in the sequence"). These holes are not an indication of queue/buffer overflow (or congestion) on the routers, and the receivers should not react accordingly.

We also make the important observation that because the sending rate is not controlled by the receivers' congestion windows, receivers with an available bandwidth higher than the transmission rate will not see congestion (or packet loss). Their window will continue to grow. This growth in the window essentially causes a large number of pending requests to be queued at the upstream router. These pending requests will allow a large number of packets to be delivered in a burst towards the receiver (e.g., when the provider decides to increase the transmit rate) This large build up of a pending requests queue at the bottleneck router (with consequent increase in the feedback delay) causes that router to indicate congestion (e.g., REM) for the large burst of packets. This causes the receiver to react and reduce the window. This form of "bang-bang" effect has the potential to under-utilize the network and/or cause unfairness between flows for a subsequent interval. While an excessively large window might cause inefficiency and unfairness, a very small window might also lead to inefficiency because a short burst caused by transient congestion might consume all the pending requests and result in unnecessary packet loss. Because of these considerations, we develop additional mechanisms to achieve congestion control and fairness.

A. Identifying Congestion & Multiplicative Decrease

We seek for help from the network to indicate congestion, using Random Early Marking (REM [26]). A simple AIMD mechanism (additive increase on receiving data and multiplicative decrease on seeing marks) works acceptably on the receivers that have available bandwidth below the sending rate of the provider. The introduction of REM into NDN can also simplify existing congestion control solutions and help avoid any unnecessary bufferbloat [27].

B. Handling Additive Increase

The optimal receiver window size can vary depending on the receiver, and can vary over time. We propose a dynamic

mechanism that observes and reacts to the instantaneous minimum pending request count (MPR) on the path, with a goal of keeping the window size close to the optimal value.

The network piggybacks the instantaneous MPR value with data packets forwarded to receivers. Note that different receivers can see different MPR values even for the same data packet since MPR is set at the branching point (at a router having other outgoing faces that can receive faster than this particular receiver). A receiver observes the MPRs over a window so as to:*1*) get the smallest value (thus avoiding excessive pending requests in the path) in that window, *2*) ensure the minimum MPR (MMPR) is larger than a small value (to avoid all pending requests being consumed by a future short burst), and *3*) smooth out variations. To simplify the calculation, our approach decides on the window increase based only on MMPR. It stops increasing the window when MMPR is larger than a configured value X and only increases window when it can rule out the possibility of reacting to a transient (based on the state machine depicted in Fig. 8, see [17] for the detailed algorithm).

We use an example to trace the MMPR and the window increase decisions to aid in understanding how the scheme works (the value of X in the state machine is set to 5 for this example). The MPR, window size and the state transitions in Fig. 7 are shown for a receiver with $2Mbps$ available bandwidth. Every cross represents a decision made at the end of a window. The MMPR grows in accordance with the window size after $1sec$. After 4 windows (decisions), the MMPR reaches 5 and the receiver stops increasing the window and MMPR stays around 5. Although the MMPR becomes lower than 5 for several windows at around $4sec$ (due to transient changes) the window size is not increased at the receiver. When there is new contention at the bottleneck link (additional competing flow) occurring between $10sec$ and $25sec$, MMPR drops to 0 and the receiver then begins to increase the window. When the receiver receives marked packets (REM) at $12sec$, it uses AIMD to respond to the congestion introduced by the competing flow. We show how this mechanism aids in achieving fairness in §VI-A2 and §VI-A3.

VI. Evaluation

We first present evaluations based on our prototype implementation. We then present simulation results for two applications (video streaming and large data file delivery) on both a small custom topology (shown in Fig. 10) and a large RocketFuel Topology [28] with heterogeneous receiver bandwiths in our custom simulation environment (used in previous work [29], [30]).

Schemes compared: *1*) An ICP solution enhanced to use REM (so as to provide a reasonable comparison with SAID). *2*) pgmcc [1], a window-based congestion control protocol which aligns the data provider transmission rate to the slowest receiver.

A. Emulation of ANP Delivery

We implemented the core components: i) ANP Delivery (described in §IV-A); and ii) congestion control (described in §V)

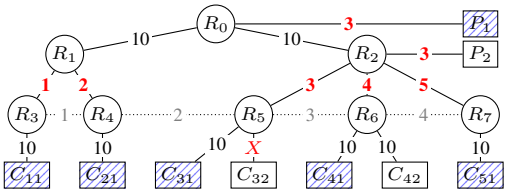

Fig. 10: Dissemination tree topology (bandwidth in $Mbps$).

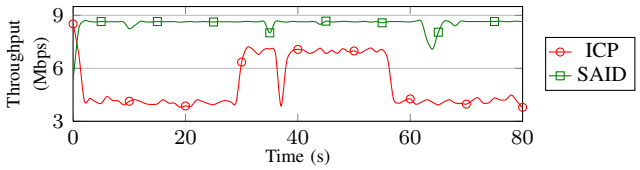

Fig. 11: Aggregate throughput in 4-consumer scenario.

in Linux[3]. The modified forwarding engine is implemented in user space and packets are encapsulated in UDP, as with CCNx.

Testbed and topology: Our testbed consists of SAID enabled routers, each running on a single machine (Ubuntu 14.0, CCNx 0.4). All the end-hosts are run on a single machine to ease the collection of results (*e.g.*, for clock synchronization reasons). The baseline topology (illustrated in Fig. 10) consists of 8 routers ($R_0 \rightarrow R_7$), 7 consumers ($C_{<x><y>}$) and 2 providers (P_1, P_2). Per link latency is 2ms and the numbers in the figure represent link bandwidth in Mbps. The bottleneck bandwidth for the consumers are marked in bold red.

1) Efficiency of ANP Delivery Model

Fig. 11 illustrates the aggregate throughput achieved by SAID and ICP in the presence of 4 consumers. pgmcc aligns to the slowest subscriber, and its results are easily obtained without needing simulations[4]. The experimental setup was used in §II-B with consumers C_{11}, C_{21}, C_{31} and C_{41}, having bottleneck bandwidths of 1, 2, 3 and 4 Mbps respectively and the provider sending at 3Mbps. SAID is able to achieve an aggregate rate close to the maximum achievable throughput (1Mbps (C_{11}) + 2Mbps (C_{21}) + 3Mbps (C_{31}) + 3Mbps (C_{41})) of 9Mbps, while ICP's is around 4Mbps. With ICP, receivers go out-of-sync and compete with one another on the link between P_1 and R_0. pgmcc can achieve approx. 4Mbps (1Mbps*4) throughput.

2) Fairness in the presence of competition

The experimental setup is P_1 with 5 consumers (named $C_{<x>1}$ in Fig. 11) and P_2 has 2 consumers (C_{32}, C_{42}). To evaluate SAID's fairness when there is a competing flow from another provider (P_2), we configure P_2 to start sending packets after approximately 10s, with a sending rate of 3Mbps (*i.e.*, as fast as it can) and study its influence in the network under an extreme case. C_{31} and C_{32} are competing for the bottleneck bandwidth between R_2 and R_5 (3Mbps) and C_{41} and C_{42} are competing for the bottleneck bandwidth between

[3]The implementation consists of roughly 1,500 lines of code and we make it publicly available in Github at https://github.com/SAIDProtocol/SAIDImplementation.git

[4]We implemented pgmcc in the simulation environment, but not the prototype environment

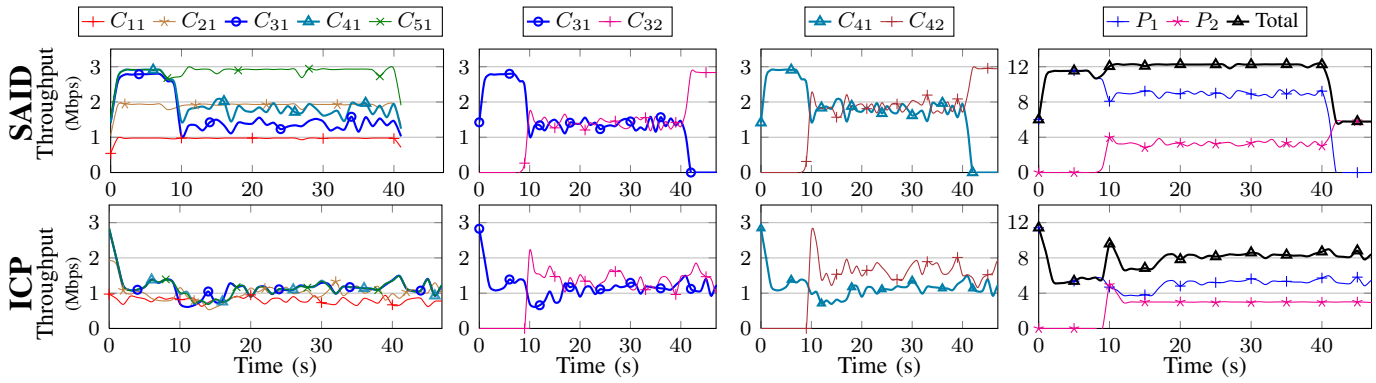

(a) P_1's consumers (b) C_{31} *vs.* C_{32} (c) C_{41} *vs.* C_{42} (d) Aggregate

Fig. 9: Emulation result in competition scenario (using topology in Fig. 10, X=10$Mbps$).

R_2 and R_6 (4Mbps). The bandwidth between R_5 and C_{32}, *i.e.* X, is set to 10Mbps so that the link between R_2 and R_5 is the bottleneck link.

Fig. 9 illustrates the throughput achieved at each receiver. As in Fig. 11, we observe in Fig. 9a that in SAID, each of P_1's receiver is able to receive at the bottleneck capacity until 10s. Once the competing flow arrives, the receivers in SAID are able to receive at their *(statistical) fair share* of the bottleneck bandwidth. In ICP, upon the arrival of a competing flow, the throughput achieved by each receiver drops significantly, from an average of 4Mbps to just 1Mbps.

Fig. 9b and Fig. 9c show that in SAID, both C_{31} and C_{41} are able to receive at their individual fair share (1.5Mbps and 2Mbps respectively) in the presence of competition from P_2's flow. On the other hand, with ICP, the consumer's of both P_1 and P_2 are not able to fully utilize their fair share of the bottleneck bandwidth, resulting in an under-utilization of the bottleneck bandwidths. This is due to the fact that both C_{31} and C_{41} are in fact competing with other consumer's from P_1 in the P_1 to R_0 link and C_{32} and C_{42} are competing with each other in the P_2 to R_0 link, due to the out-of-sync phenomena. The peak at the start of the P_2 flows in ICP, illustrate that for a short period, flows to C_{32} and C_{42} are in sync.

Finally, in Fig. 9d, the aggregate throughput achieved by P_1's and P_2's consumers is close to the maximum achievable throughput with SAID. The total throughput of the network in SAID, up to 10s is close to the ideal 12Mbps. After 10s is even closer to the ideal 13Mbps. With ICP however, it is approximately 5Mbps before 10s and 8Mbps afterwards. In pgmcc, in the absence of a competitor (*i.e.*, before 10s), P_1 aligns to C_{11} (1mbps) and therefore the total throughput is 5Mbps (since all its 5 consumers receive at 1Mbps each). In the presence of a competitor P_2 (after 10s) who is sending at 2Mbps (aligned to C_{32}), the total throughput will increase to 9Mbps.

3) Benefit of receiver controlled mechanism

We now modify our experimental setup to a simple dumb-bell topology used in [25] by changing the link bandwidth between R_5 and C_{32}, *i.e.*, X to 0.3Mbps. Since C_{32} can receive only at 0.3Mbps from P_2, C_{31} is able to make use of the spare capacity of its bottleneck link between R_2 and R_5 and receive

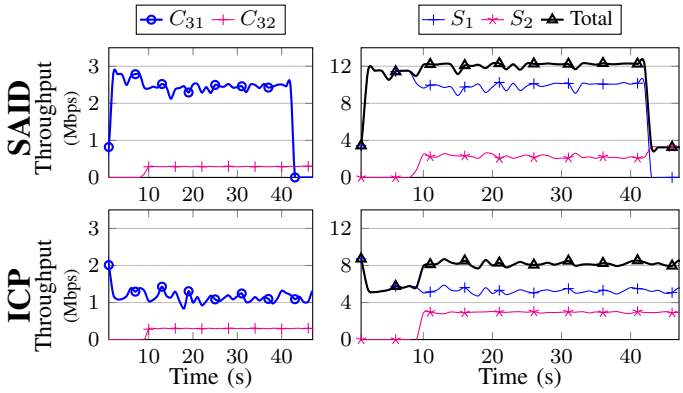

(a) C_{31} *vs.* C_{32} (b) Aggregate

Fig. 12: Emulation result in competition scenario (using topology in Fig. 10, X=300$kbps$).

at approximately 2.7Mbps (See Fig. 12a), thereby making optimal use of that link (total usage is 2.7Mbps + 0.3Mbps). On the other hand, a flow based fair queuing approach would only achieve 1.8Mbps (1.5Mbps(C_{31})+ 0.3Mbps (C_{32})), since the link between R_2 and R_5 would be shared equally between flows from P_1 and P_2, only for P_2's packets to be dropped at R_5. This result highlights the benefit of using a receiver-driven approach for SAID. In the case of ICP, C_{31} is only able to receive close to an equal share of the R_2 and R_5 link bandwidth (¡1.5Mbps). ICP gets approx 1.8Mbps aggregate throughput on that link. The total throughout for SAID (12 Mbps) and ICP (8 Mbps) are shown in Fig. 12a. For pgmcc, in the presence of a competitor, the total achievable throughput is 5.6Mbps (1Mbps * 5(C_{11}-C_{51}) + 0.3Mbps *2(C_{32},C_{42})).

4) Multiple Bottlenecks: Fairness with SAID

We illustrate the fairness of SAID's receiver-driven congestion control mechanism with a configuration having multiple bottlenecks (Generic Fairness Configuration (GFC [31])) so that different flows have different fair rates. The topology and the link capacities (in $Mbps$) is shown in Fig. 13. There are 6 flow groups ($A-F$) each with different # of flows (in the bracket after the group name). The ideal fair throughput of the flows are listed in Table II (in the "Fair" column).

Since SAID uses a receiver-side congestion window to

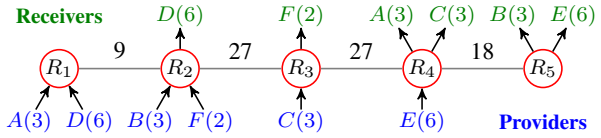

Fig. 13: Generic Fairness Configuration (GFC).

TABLE II: Fairness Result in GFC.

Group	Flows	Hops	Fair	SAID		ICP	
				Avg.	Avg./Fair	Avg.	Avg./Fair
A	3	5	1	0.7347	0.7347	0.7245	0.7245
B	3	5	2	1.2352	0.6176	1.2875	0.6437
C	3	3	6	6.5399	1.0900	6.4909	1.0818
D	6	3	1	1.1284	1.1284	1.1363	1.1363
E	6	3	2	2.3653	1.1826	2.3522	1.1761
F	2	3	9	9.0665	1.0047	9.0427	1.0047

TABLE III: Stall time (s) in streaming demo (video length=$40s$).

	C_1	C_2	C_3	C_4	Rep.
Baseline	83.384	22.507	2.461	0.886	–%
ICP	90.530	33.965	33.821	33.820	–%
pgmcc	84.804	84.770	84.768	84.767	0.00%
SAID-F	83.821	40.062	39.569	1.131	12.44%
SAID-S	83.541	22.754	4.010	1.123	21.91%
SAID	44.304	1.271	1.151	1.131	12.44%

control the receive rate, a feedback loop is formed between the receiver and a branching point. Thus the change in the congestion window of requests does not affect the other receivers or the data provider. Every "flow" in GFC can be seen as part of a dissemination tree (path from provider to receiver). The data providers transmit data faster than any of the receivers' rate (at $27Mbps$). After ignoring start up effects, we determine the average throughput of each of the flows in the subsequent $15sec$ period.

The average throughput and the fairness ratio ($\frac{avg.}{fair}$) for each flow with the two protocols are listed in Table II. Just like other feedback window-based congestion control mechanisms, SAID slightly benefits the flows with shorter RTT ($C-F$). Similar results are observed for ICP. In general, SAID performs similarly to ICP (and all the other window-based congestion control mechanisms) from a fairness standpoint. We believe SAID would also be fair to protocols using AIMD on *each* path of the dissemination tree separately.

B. Simulation of different applications

1) Video Streaming

We consider a streaming video application with a playout rate of $3Mbps$. We evaluate the effectiveness of the approaches using the *Stall time*, which reflects the impact on user experience. No stalls occur when there are no holes in the sequence in the play out buffer for the next $1s$ of video playback. For the baseline, we run the simulation 4 times with a single receiver requesting the video. SAID-F is when repair is by the provider at the end of the flow (similar to "download-and-play"). SAID-S is repair by the provider while streaming the video. SAID is the peer-assisted repair solution we propose here. These solutions are compared on the tree topology (Fig. 10, with only C_{11}-C_{41} activated) as well as ICP and pgmcc.

Table III shows the stall time and the repair ratio ($\frac{\text{\# of pkts via repair}}{\text{\# of total packets}}$, 'Rep' in Table) for a $40s$ video. In ICP, when all receivers request the video simultaneously, they go *out-of-sync* soon thereafter. The stall time becomes larger than the baseline especially for faster receivers. With pgmcc, since the provider has to align with the slowest receiver (C_1), the stall time for the rest of the receivers is up to $80s$ larger compared to the baseline. The user experience for the faster receivers therefore deteriorates considerably.

Although SAID-F achieves a relatively low repair rate, the stall time for C_1-C_3 remains high, because the repair is performed after the flow finishes. C_2 and C_3 benefit from the repair during streaming in SAID-S. But this benefit comes at the cost of higher network load. The repair ratio of SAID-S is higher than SAID-F ($\sim$22% *vs.* $\sim$12.5%) since the retransmission has to go through the bottleneck link and affects the primary "any packet" stream. But the retransmission rate from the provider of SAID-S is still lower than ICP ($\sim$1.9 *vs.* $\sim$3.4, not shown) which means SAID-S consumes less network and provider resources. SAID (our proposal) however is superior as it is able to utilize the extra bandwidth between end-hosts. The repair does not affect the multicast session and the slower consumers can get twice the bandwidth compared to SAID-F and SAID-S. Despite the repair, the stall time for the slower receivers (C_1-C_3) is much smaller than with the other solutions, even though they are playing the video at the same rate of $3Mbps$. We varied the video length and observed that the stall time grows proportionally with the video length but the pattern in Table III still holds.

2) File Content Delivery

We evaluated the complete SAID solution with the Rocket-Fuel topology (Fig. 14a). We randomly place 20-200 receivers on the 79 core routers in the topology. Since RocketFuel does not have bandwidth information, we assign available bandwidths in the range of 1-10$Mbps$ for each link. The result of a trace with 100 flows using ICP, SAID and pgmcc is in Fig. 14.

By decoupling reliability from congestion control, SAID has lower network load compared to ICP (Fig. 14b), especially as the number of receivers increases, by up to 46% with 200 receivers. Since $\sim$60% of the data is delivered at the first attempt, SAID has a much lower average # of transmissions of each packet from the provider and average flow completion time compared to ICP (Fig. 14d). But SAID consumes more network bandwidth ($\sim$10%) compared to pgmcc since it aligns to a faster receiver and uses repairs subsequently. For the 10% additional network load, SAID achieves lower average completion time (by $\sim$65%, Fig. 14c).

VII. Related Work

Here we focus on the related work that are not mentioned earlier.

Many reliable multicast protocols have been proposed to enable large scale reliable data dissemination. To maximize the utility of multicast, cyclic- and scheduled-multicasts [32], [33] have been proposed to benefit the consumers that are not starting at the same time. These solutions can be broadly

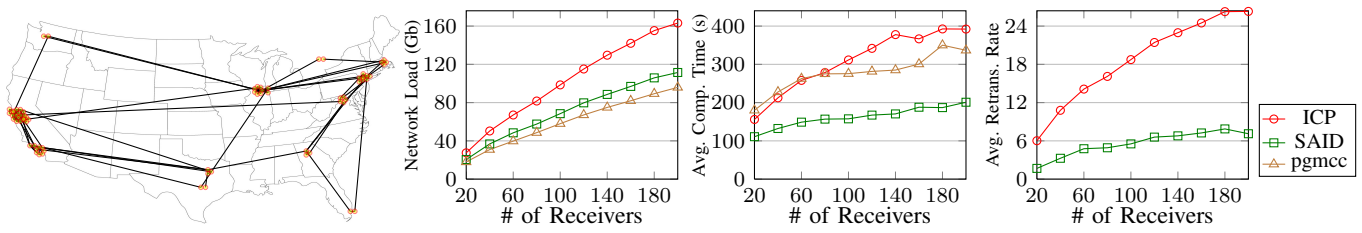

(a) RocketFuel topology (Exodus, AS-3967) (b) Aggregate network load (c) Average Completion time (d) Retransmission rate

Fig. 14: Simulation result for file content delivery application.

classified into two categories based on the chosen repair mechanism: provider repair or peer-assisted repair.

The clients in provider-repair mechanisms [34]–[39] send NAKs to the provider (or the whole group to suppress duplicate NAKs) and the provider retransmits the missing packets specified in the NAKs. The data provider in such solutions has to align the sending rate to the slowest receiver eventually.

Peer-assisted approaches like [15], [21]–[23], [40], [41] group receivers in a hierarchical structure and the receivers ACK to the upper level in the tree so that the ACKs can be aggregated. By introducing such a relationship among receivers, these solutions allow receivers to perform local repair and therefore, the provider can align the sending rate to the majority of (or the fastest) receivers based on the ACK strategy. Unfortunately, in these proposals, the subscribers have to exchange information in a peer-to-peer manner in the IP network to perform repair as well as send ACKs. According to [24], these solutions face the problem of **privacy** and **trust**. Namely, the receivers have to reveal their identities (IP addresses) to the peers, and thereby trust them as there is no guarantee of data integrity during peer-repair.

Layered multicasts are also proposed to deal with heterogeneous consumers. In [42], [43], the provider creates different multicast groups that transmit different resolutions of the data. The receivers can select appropriate groups according to their link capacity. These solutions are applicable to select applications. But having a reliable multicast capability for a single rate stream is still fundamental for broad-based use across all kinds of applications. Therefore, the layered multicast solutions can be considered as orthogonal to the single-rate reliable multicast.

Many congestion-control mechanisms have been proposed for ICN like [5]–[9]. They all face the out-of-sync issue since they rely on similar TCP-like mechanism. Other works such as [29], [44] try to achieve efficient large scale data dissemination via pub/sub. But they lack an efficient mechanism to ensure reliability and avoid congestion collapse in the network.

VIII. Conclusion

Designing a congestion control mechanism for information delivery to large numbers of receivers is difficult, particularly with heterogeneous receiver bandwidths. It continues to be a challenge also with Information Centric Networks. Through emulation and an analytical model, we showed that heterogeneous receivers will get out-of-sync with existing receiver-driven, in-sequence congestion control mechanisms. To overcome the resulting inefficiencies, we proposed SAID. SAID allows receivers to request "any-next" packet, instead of the "next in-sequence" and thus delivers more packets on the first attempt. For missing packets, other receivers can provide "repair" packets even if the in-network caches no longer hold the content. Privacy and trust is maintained during the repair phase.

Our evaluations show that SAID achieves efficiency and fairness on each path between the information provider and receivers. From a large scale simulation, SAID can reduce aggregate network load (by up to $\sim$46%) and transmission completion times (by more than 50%) compared to existing congestion control mechanisms. SAID also reduces completion time by $\sim$40% while only increasing network load by $\sim$10% compared to pgmcc, which aligns the sending rate to the slowest receiver. Based on measurements on a prototype, SAID outperforms ICP, getting 50% higher aggregate throughput and almost twice the throughput of pgmcc. With the efficient repair of SAID, streaming applications have a much smaller stall time compared to the other mechanisms.

Acknowledgment

The research leading to these results has received funding from the EU-JAPAN initiative by the EC Seventh Framework Programme (FP7/2007-2013) Grant Agreement No. 608518 (GreenICN), NICT under Contract No. 167, the US National Science Foundation under Grant No. CNS-1455815, and the Volkswagen Foundation Project "Simulation Science Center". The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the GreenICN project, the Simulation Science Center project, the European Commission, NSF, or NICT.

References

[1] L. Rizzo, "pgmcc: a tcp-friendly single-rate multicast congestion control scheme," in *SIGCOMM*, 2000.

[2] M. Castro, P. Druschel, A.-M. Kermarrec, and A. I. Rowstron, "SCRIBE: A large-scale and decentralized application-level multicast infrastructure," *JSAC*, pp. 1489–1499, 2002.

[3] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking Named Content," in *CoNEXT*, 2009.

[4] L. Zhang, D. Estrin, J. Burke, V. Jacobson, and J. Thornton, "Named Data Networking (NDN) Project," PARC, Tech. Report NDN-0001, 2010.

[5] G. Carofiglio, M. Gallo, and L. Muscariello, "ICP: Design and evaluation of an interest control protocol for content-centric networking," in *ICN*, 2012.

[6] ——, "Joint hop-by-hop and receiver-driven interest control protocol for content-centric networks," in *ICN*, 2012.
[7] S. Salsano, A. Detti, M. Cancellieri, M. Pomposini, and N. Blefari-Melazzi, "Transport-layer issues in information centric networks," in *ICN*, 2012.
[8] A. Somaya, N. Pekka, E. Lars, O. Jorg, and W. Walter, "Contug: A receiver-driven transport protocol for content-centric networks," in *ICNP (poster session)*, 2010.
[9] L. Saino, C. Cocora, and G. Pavlou, "Cctcp: A scalable receiver-driven congestion control protocol for content centric networking," in *ICC*, 2013.
[10] V. Jacobson, "Congestion avoidance and control," in *SIGCOMM*, 1988.
[11] K. Ramakrishnan and R. Jain, "A binary feedback scheme for congestion avoidance in computer networks," *TOCS*, pp. 158–181, 1990.
[12] S. Oueslati, J. Roberts, and N. Sbihi, "Flow-aware traffic control for a content-centric network," in *INFOCOM*, 2012.
[13] S. Arianfar, P. Nikander, and J. Ott, "On content-centric router design and implications," in *ReArch*, 2010.
[14] V. Jacobson, D. K. Smetters, N. H. Briggs, M. F. Plass, P. Stewar, J. D. Thornton, and R. L. Braynard, "VoCCN: voice-over content-centric networks," in *ReArch*, 2009.
[15] S. Floyd, V. Jacobson, C.-G. Liu, S. McCanne, and L. Zhang, "A reliable multicast framework for light-weight sessions and application level framing," *TON*, pp. 784–803, 1997.
[16] X. Zhang, J. Liu, B. Li, and T.-S. P. Yum, "CoolStreaming/DONet: A Data-driven Overlay Network for Peer-to-Peer Live Media Streaming," in *INFOCOM*, 2005.
[17] J. Chen, M. Arumaithurai, X. Fu, and K. K. Ramakrishnan, "SAID: A scalable and adaptive information dissemination protocol in icn," Tech. Rep., 2014. [Online]. Available: https://www.net.informatik.uni-goettingen.de/publications/1912/SAID.pdf
[18] J. He and J. Rexford, "Toward Internet-wide Multipath Routing," *Network*, pp. 16–21, 2008.
[19] X. Liu and L. Xiao, "A Survey of Multihoming Technology in Stub Networks: Current Research and Open Issues," *Network*, pp. 32–40, 2007.
[20] P. Key, L. Massoulié, and D. Towsley, "Path selection and multipath congestion control," in *INFOCOM*, 2007.
[21] Y. Chawathe, S. McCanne, and E. A. Brewer, "Rmx: Reliable multicast for heterogeneous networks," in *INFOCOM*, 2000.
[22] S. Ratnasamy, A. Ermolinskiy, and S. Shenker, "Revisiting ip multicast," in *SIGCOMM CCR*, vol. 36, no. 4, 2006, pp. 15–26.
[23] A. Malekpour, A. Carzaniga, F. Pedone, and G. Toffetti Carughi, "End-to-end reliability for best-effort content-based publish/subscribe networks," in *DEBS*, 2011.
[24] N. Duffield, M. Grossglauser, and K. Ramakrishnan, "Distrust and privacy: Axioms for multicast congestion control," in *NOSSDAV*, 1999.
[25] S. Floyd and K. Fall, "Promoting the use of end-to-end congestion control in the internet," *TON*, pp. 458–472, 1999.
[26] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *TON*, vol. 1, no. 4, pp. 397–413, 1993.
[27] J. Gettys and K. Nichols, "Bufferbloat: dark buffers in the internet," *Communications of the ACM*, pp. 57–65, 2012.
[28] R. Mahajan, N. Spring, D. Wetherall, and T. Anderson, "Inferring Link Weights using End-to-End Measurements," in *IMW*, 2002.
[29] J. Chen, M. Arumaithurai, L. Jiao, X. Fu, and K. K. Ramakrishnan, "COPSS: An Efficient Content Oriented Pub/Sub System," in *ANCS*, 2011.
[30] J. Chen, M. Arumaithurai, X. Fu, and K. K. Ramakrishnan, "Coexist: Integrating Content Oriented Publish/Subscribe Systems with IP," in *ANCS*, 2012.
[31] R. J. Simcoe, "Test configurations for fairness and other tests," in *ATM Forum/94-0557*, 1994.
[32] K. C. Almeroth, M. H. Ammar, and Z. Fei, "Scalable delivery of web pages using cyclic best-effort multicast," in *INFOCOM*, 1998.
[33] V. Aggarwal, R. Caldebank, V. Gopalakrishnan, R. Jana, K. Ramakrishnan, and F. Yu, "The effectiveness of intelligent scheduling for multicast video-on-demand," in *Multimedia*, 2009.
[34] D. Towsley, "An analysis of a point-to-multipoint channel using a go-back-n error control protocol," *Communications, IEEE Transactions on*, pp. 282–285, 1985.
[35] K. Sabnani and M. Schwartz, "Multidestination protocols for satellite broadcast channels," *Communications, IEEE Transactions on*, pp. 232–240, 1985.
[36] A. Erramilli and R. Singh, "A reliable and efficient multicast protocols for broadband broadcast networks," in *SIGCOMM*, 1987.
[37] S. Ramakrishnan and B. Jain, "A negative acknowledgement with periodic polling protocol for multicast over lans," in *INFOCOM*, 1987.
[38] S. Paul, K. K. Sabnani, and D. M. Kristol, "Multicast transport protocols for high speed networks," in *ICNP*, 1994.
[39] B. Adamson, C. Bormann, M. Handley, and J. Macker, "Nack-oriented reliable multicast (norm) transport protocol," *IETF RFC*, vol. 5740, 2009.
[40] J. C. Lin and S. Paul, "Reliable multicast transport protocol," in *INFOCOM*, 1996.
[41] D. M. Chiu, S. Hurst, M. Kadansky, and J. Wesley, "Tram: A tree-based reliable multicast protocol," 1998.
[42] S. McCanne, V. Jacobson, and M. Vetterli, "Receiver-driven layered multicast," in *SIGCOMM*, 1996.
[43] L. Vicisano, J. Crowcroft, and L. Rizzo, "Tcp-like congestion control for layered multicast data transfer," in *INFOCOM*, 1998.
[44] N. Fotiou, P. Nikander, D. Trossen, and G. C. Polyzos, "Developing information networking further: From psirp to pursuit," in *BROADNETS*, 2012, pp. 1–13.